arXiv:1407.6730v3 [cs.DS] 4 Aug 2014

New routing techniques and their applications

Liam Roditty* Bar Ilan University

Roei Tov† Bar Ilan University

August 5, 2014

Abstract

Let $G = (V, E)$ be an undirected graph with n vertices and m edges. We present two new routing techniques. Roughly speaking, given a partition $\mathcal{U} = \{U_1, \ldots, U_q\}$ of V into q sets each of size $\tilde{O}(n/q)$, our first technique routes a message between vertices of $U \in \mathcal{U}$ on a $(1+\varepsilon)$-stretch path with routing tables of size $\tilde{O}(\frac{1}{\varepsilon}(n/q) + q)$. Given a partition $\mathcal{W} = \{W_1, \ldots, W_q\}$ of a set $W \subseteq V$ into q sets each of size $\tilde{O}(|W|/q)$ and assuming that the sets of $\mathcal{U}$ satisfy a certain hitting set property with respect to vertex vicinities, our second technique routes a message, for every $i \in \{1, \ldots, q\}$, from any source in U_i to any destination in W_i on a $(1+\varepsilon)$-stretch path with routing tables of size $\tilde{O}(\frac{1}{\varepsilon}(|W|/q) + q)$. Using these techniques we obtain the following new routing schemes:

- A routing scheme for unweighted graphs that uses $\tilde{O}(\frac{1}{\varepsilon}n^{2/3})$ space at each vertex and $\tilde{O}(1/\varepsilon)$-bit headers, to route a message between any pair of vertices $u, v \in V$ on a $(2+\varepsilon, 1)$-stretch path, i.e., a path of length at most $(2+\varepsilon) \cdot d + 1$, where d is the distance between u and v. This should be compared to the $(2,1)$-stretch and $\tilde{O}(n^{5/3})$ space distance oracle of Pătraşcu and Roditty [FOCS'10 and SIAM J. Comput. 2014] and to the $(2,1)$-stretch routing scheme of Abraham and Gavoille [DISC'11] that uses $\tilde{O}(n^{3/4})$ space at each vertex. It follows from Pătraşcu, Thorup and Roditty [FOCS'12] that a 2-stretch routing scheme with $\tilde{O}(m^{2/3})$ space at each vertex is optimal, assuming a hardness conjecture on set intersection holds.
- A routing scheme for weighted graphs with normalized diameter D, that uses $\tilde{O}(\frac{1}{\varepsilon}n^{1/3}\log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon}\log D)$-bit headers, to route a message between any pair of vertices on a $(5+\varepsilon)$-stretch path. This should be compared to the 5-stretch and $\tilde{O}(n^{4/3})$ space distance oracle of Thorup and Zwick [STOC'01 and J. ACM. 2005] and to the 7-stretch routing scheme of Thorup and Zwick [SPAA'01] that uses $\tilde{O}(n^{1/3})$ space at each vertex. Since a 5-stretch routing scheme must use tables of $\Omega(n^{1/3})$ space our result is almost tight.
- For an integer $\ell > 1$, a routing scheme for unweighted graphs that uses $\tilde{O}(\ell\frac{1}{\varepsilon}n^{\ell/(2\ell\pm1)})$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon})$-bit headers, to route a message between any pair of vertices on a $(3 \pm 2/\ell + \varepsilon, 2)$-stretch path.
- A routing scheme for weighted graphs, that uses $\tilde{O}(\frac{1}{\varepsilon}n^{1/k}\log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon}\log D)$-bit headers, to route a message between any pair of vertices on a $(4k-7+\varepsilon)$-stretch path.

*liam.roditty@biu.ac.il

†roei81@gmail.com

1 Introduction

Graph spanners, *distance oracles* and *compact routing schemes* are fundamental notions in graph theory, data structures and distributed algorithms, respectively, that deal with the natural tradeoff between space and accuracy.

Peleg and Ullman [17] and Peleg and Schäffer [16] introduced the notion of graph spanners. Let $G = (V, E)$ be an undirected graph. For $u, v \in V$, let $d(u, v)$ be the length of a shortest path between u and v. A path between u and v is of (α, β)-stretch if its length is at most $\alpha \cdot d(u, v) + \beta$. For $(\alpha, 0)$-stretch we write α-stretch. A graph $H = (V, E')$, where $E' \subseteq E$, is an (α, β)-stretch spanner of G, if and only if, for every $u, v \in V$ there is an (α, β)-stretch path between u and v in H. It is known how to efficiently construct a $(2k-1)$-spanner of size $\tilde{O}(n^{1+1/k})$ [5, 9], and this size-stretch tradeoff is tight assuming the *girth* conjecture of Erdős [14] holds.

Thorup and Zwick [22] introduced the notion of distance oracles. They showed that it is possible to preprocess a weighted undirected graph in $O(mn^{1/k})$ expected time and create a data structure of size $O(n^{1+1/k})$ that can answer distance queries with a $(2k-1)$-stretch between any two vertices in $O(k)$ time. For $k = 1$ this gives the trivial solution of exact distances with $O(n^2)$ space. For $k = 2$, this gives a 3-stretch distance oracle with $O(n^{3/2})$ space. Pătraşcu and Roditty [19] showed that there is another distance oracle between these two solutions. They presented a $(2, 1)$-stretch distance oracle with $\tilde{O}(n^{5/3})$ space. Pătraşcu, Thorup and Roditty [20] generalized this and showed that for every $\ell \geq 1$, there is a $(3 - 2/\ell)$-stretch distance oracle with $\tilde{O}(\ell m^{1+\ell/(2\ell-1)})$.

Peleg and Upfal [18] introduced the notion of compact routing schemes. Awerbuch, Bar-Noy, Linial and Peleg [7] were the first to distinguish between *labeled* routing schemes and *name independent* routing schemes. In compact routing schemes there is a preprocessing phase in which a centralized algorithm computes routing tables. In *labeled* routing schemes a short label is assigned to each vertex. After the preprocessing phase ends, messages with an additional short header can be routed between the vertices of the graph in a distributed manner. More specifically, when a message to destination v reaches a vertex $u \neq v$, the routing scheme algorithm executed at u decides locally, based on the routing table of u, the message header and the label of v on which link to forward the message. For an integer $k > 1$, Thorup and Zwick [21] presented a $(4k - 5)$-stretch labeled compact routing scheme with routing tables of size $\tilde{O}(n^{1/k})$-bit at each vertex, $o(k \log^2 n)$-bit labels and $o(\log^2 n)$-bit headers.

In a sense finding a good spanner is easier than finding a good distance oracle since a distance oracle has to support efficient distance queries between any pair of vertices. Obtaining a routing scheme is inherently harder than obtaining a distance oracle since decisions are made locally at a vertex based on a fraction of the whole information, while in distance oracles the whole information is always available.

Therefore, a general open problem is given a (α, β)-stretch S-space distance oracle can we also obtain a (α, β)-stretch routing scheme with $O(S/n)$-space routing tables? Both spanners and distance oracles have the same general tradeoff of $(2k - 1)$-stretch and $O(n^{1+1/k})$-space which is optimal under the girth conjecture. An important challenge is to obtain also a routing scheme with $(2k - 1)$-stretch and $\tilde{O}(n^{1/k})$-space routing tables.

Recently, Chechik [10] made an important breakthrough and improved the stretch of the general routing scheme of Thorup and Zwick [21] from $4k - 5$ to $((4 - \alpha)k - \beta)$-stretch, for some constants α and β. The routing tables in her scheme are of size $O(n^{1/k} \log D)$, where D is the normalized diameter of the graph.

Routing schemes, and in particular small stretch routing schemes, have been extensively studied over the last three decades. Cowen [11] obtained a 3-stretch labeled routing scheme with routing tables of $\tilde{O}(n^{2/3})$ space. Eilam *et al.* [13] obtained a 5-stretch labeled routing scheme with routing tables of $\tilde{O}(n^{1/2})$ space. Arias *et al.* [6] obtained a 5-stretch name independent routing scheme with routing tables of $\tilde{O}(n^{1/2})$ space.

Thorup and Zwick [21] improved the result of Cowen [11] and obtained a 3-stretch labeled routing scheme with routing tables of $\tilde{O}(n^{1/2})$ space. Finally, Abraham *et al.* [3] obtained a 3-stretch name independent routing scheme with routing tables of $\tilde{O}(n^{1/2})$ space.

Remarkably, these last two routing schemes are the only optimal routing schemes for general undirected graphs amongst all the possible optimal stretch/space combinations. Even the new routing scheme of Chechik [10] is better than the $(4k-5)$-stretch routing scheme only for $k \geq 4$. In particular, for $k = 4$ the stretch is approximately 10.52 instead of 11 as in the routing scheme of [21]. Therefore, we are evidence to a phenomena in which settling on any stretch factor from the range $(3, 7)$ cannot help in reducing the routing table size below the $\sqrt{n}$ barrier.

In this paper we break this long standing barrier and present a $(5+\varepsilon)$-stretch routing scheme that uses $\tilde{O}(\frac{1}{\varepsilon}n^{1/3}\log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon}\log D)$-bit headers. The label of each vertex is of $O(\log n)$ size. This almost matches the optimal stretch/space combination of 5-stretch and $O(n^{1/3})$-space routing tables.

Following the $(2,1)$-stretch distance oracle of [19] for unweighted graphs, Abraham and Gavoille [1] presented a $(2,1)$-stretch routing scheme with routing tables of size $\tilde{O}(n^{3/4})$ and in general a $(4k-6,1)$-stretch $\tilde{O}(n^{3/(3k-2)})$-space routing tables. In this paper we almost match the $(2,1)$-stretch distance oracle of [19] and present a $(2+\varepsilon,1)$-stretch routing scheme with routing tables of $\tilde{O}(\frac{1}{\varepsilon}n^{2/3})$ size at each vertex and $\tilde{O}(1/\varepsilon)$-bit headers. In Table 1 we compare our new results for small stretch factors with the previous best bounds.

We obtain our new results by extending a fundamental idea for routing between nearby vertices. In many of the existing routing schemes every vertex $u \in V$ stores the first edge on a shortest path to every vertex in $B(u,\ell)$, the set of ℓ closest vertices of u. A message is routed from u to $v \in B(u,\ell)$ using the edge that u stored for v. It is easy to show that $v \in B(w,\ell)$ for every w that is on a shortest path between u and v. Thus, it is possible to route a message from u to a specific vertex $v \in B(u,\ell)$ on a shortest path by storing at u $O(\log n)$-bits of routing information dedicated for v.

At the heart of our new routing techniques is an extension of this fundamental idea. Roughly speaking, we show that if every $u \in V$ stores the first edge on a shortest path to every $v \in B(u,\ell)$, and an additional routing information on a fixed set of special vertices, it is possible to route a message from u to a specific vertex v that is arbitrarily far away from u on a $(1+\varepsilon)$-stretch path by storing at u $O(\frac{1}{\varepsilon}\log n)$-bits of routing information dedicated for v.

Using this extension we present two new routing techniques. Given a partition $\mathcal{U} = \{U_1, \ldots, U_q\}$ of V into q sets each of size $\tilde{O}(n/q)$, our first technique routes a message between any pair of vertices of $U \in \mathcal{U}$ on a $(1+\varepsilon)$-stretch path with routing tables of size $\tilde{O}(\frac{1}{\varepsilon}(n/q)+q)$. Given a partition $\mathcal{W} = \{W_1, \ldots, W_q\}$ of a set $W \subseteq V$ into q sets each of size $\tilde{O}(|W|/q)$ and assuming that $U \cap B(u, \tilde{O}(q)) \neq \emptyset$, for every $U \in \mathcal{U}$ and $u \in V$, our second technique routes a message, for every $i \in \{1, \ldots, q\}$, from any source in U_i to any destination in W_i on a $(1+\varepsilon)$-stretch path with routing tables of size $\tilde{O}(\frac{1}{\varepsilon}(|W|/q)+q)$.

Our $(5+\varepsilon)$-stretch routing scheme is obtained using our second routing technique together with the routing scheme of Thorup and Zwick [21] and a coloring method used by Abraham *et al.* [3] and Abraham and Gavoille [1]. Our $(2+\varepsilon,1)$-stretch routing scheme is obtained using our first routing technique together with ideas of Abraham and Gavoille [1] and Pătraşcu and Roditty [19]. Using our first technique it is also possible to obtain a name independent routing scheme with stretch $3+\varepsilon$ and routing tables of $\tilde{O}(\sqrt{n})$ size.

Our second technique is strong enough to obtain also the following generalizations:

- For an integer $\ell > 1$, a routing scheme for unweighted graphs that uses $\tilde{O}(\ell\frac{1}{\varepsilon}n^{\ell/(2\ell\pm1)})$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon})$-bit headers, to route a message between any pair of vertices on a $(3 \pm 2/\ell + \varepsilon, 2)$-stretch path. This almost matches the distance oracles of Pătraşcu, Thorup and Roditty [20].

Reference	Graph	Stretch	Table Size
Abraham and Gavoille [1]	Unweighted	$(2,1)$	$\tilde{O}(n^{3/4})$
Thorup and Zwick [21], Abraham *et al.* [3]	Weighted	3	$\tilde{O}(n^{1/2})$
Thorup and Zwick [21]	Weighted	7	$\tilde{O}(n^{1/3})$
Chechik [10]	Weighted	10.52	$\tilde{O}(n^{1/4}\log D)$
Theorem 10	Unweighted	$(2+\varepsilon,1)$	$\tilde{O}((1/\varepsilon)\cdot n^{2/3})$
Theorem 13	Unweighted	$(2\frac{1}{3}+\varepsilon,2)$	$\tilde{O}((1/\varepsilon)\cdot n^{3/5})$
Theorem 15	Unweighted	$(4+\varepsilon,2)$	$\tilde{O}((1/\varepsilon)\cdot n^{2/5})$
Theorem 11	Weighted	$5+\varepsilon$	$\tilde{O}((1/\varepsilon)\cdot n^{1/3}\log D)$
Theorem 16	Weighted	$9+\varepsilon$	$\tilde{O}((1/\varepsilon)\cdot n^{1/4}\log D)$

Table 1: Previously available routing schemes and our new routing schemes

- A routing scheme that uses $\tilde{O}(\frac{1}{\varepsilon}n^{1/k}\log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon}\log D)$-bit headers, to route a message between any pair of vertices on a $(4k-7+\varepsilon)$-stretch path.

Our $(4k-7+\varepsilon)$-stretch routing scheme can be used in the $((4-\alpha)k-\beta)$-stretch routing scheme of Chechik [10] instead of the $(4k-5)$-stretch routing scheme of Thorup and Zwick in order to obtain slightly better values for α and β.

The rest of this paper is organized as follows. In the next section we present some preliminaries that are needed throughout the paper. In Section 3 we present our new routing techniques. In Section 4 we present our new routing schemes for small stretch factors. In Section 5 we present our generalized routing schemes.

2 Preliminaries

Let $G=(V,E)$ be an n-vertices m-edges undirected graph. For every $u,v\in V$, let $d(u,v)$ be the length of a shortest path between u and v. Let $D=\frac{\max_{u,v} d(u,v)}{\min_{u\neq v} d(u,v)}$ be the *normalized diameter* of G. Let $B(u,\ell)$ be the ℓ closest vertices of u, breaking ties by lexicographical order of vertex names. Let $r_u(\ell)$ be the largest value for which it holds that every vertex $w\in V$ with $d(u,w)=r_u(\ell)$ is in $B(u,\ell)$. When it will be clear from the context we simply write r_u. Notice that for unweighted graphs it follows from this definition that $d(u,w)\leq r_u(\ell)+1$ for every $w\in B(u,\ell)$. The following property is a well known property of such vertex vicinities.

Property 1 ([8],[2])**.** *If $v\in B(u,\ell)$ and w is on a shortest path between u and v then $v\in B(w,\ell)$.*

The next Lemma is a direct result of the above property.

Lemma 2. *Let $G=(V,E)$ be a weighted graph and let $\ell>0$ be an integer. We can construct a routing scheme that uses routing tables of size $O(\ell)$ in each vertex and $O(\log n)$-bit header, such that a message is routed from u to any $v\in B(u,\ell)$ on a shortest path.*

Proof. Let $u\in V$. For each $v\in B(u,\ell)$ store at u the first edge on a shortest path to v. Assume $v\in B(u,\ell)$ and let (u,w) be the edge saved at u for v. When u receives a message to v, it forwards it to w. From Property 1 it follows that $v\in B(w,\ell)$, therefore, w has the next edge on a shortest path to v and the message is routed from u to v on a shortest path. □

Throughout the paper when we say that u saves $B(u,\ell)$ we mean that u can check membership and to retrieve the saved edge in constant time.

For our routing schemes we assume the standard fixed port model as described by Fraigniaud and Gavoille [15]. We use the following tree routing scheme as a building block in our routing schemes:

Lemma 3 (Tree routing [21, 15]). *Given a weighted tree there is a routing scheme that given the label of a destination vertex in the tree, routes from any source vertex in the tree to the destination on the shortest path in the tree. The storage of a node, its label size and the header size are $O(\log^2 n/\log\log n)$ bits. Each link on the path is obtained in constant time.*

The notions of bunches and clusters were used by Thorup and Zwick in the context of distance oracles [22]. Let $A \subseteq V$. Let $p_A(u) = \arg\min\{d(u,v) \mid v \in A\}$, breaking ties by lexicographical order of vertex names, and $d(u,A) = d(u,p_A(u))$.

The *bunch* of $v \in V$ is $B_A(v) = \{w \in V \mid d(w,v) < d(v,A)\}$. The *cluster* of $w \in V$ is $C_A(w) = \{v \in V \mid d(w,v) < d(v,A)\}$. Notice that $w \in B_A(v)$ if and only if $v \in C_A(w)$.

From the definition of clusters it follows that if $u \in C_A(w)$ and v is on a shortest path between u and w then $v \in C_A(w)$. Thus, there is a shortest path tree $T_{C_A}(w)$ rooted at w that spans the vertices of $C_A(w)$. Let $|A| = s$, the size of each bunch is bounded by $O(n/s)$. Thorup and Zwick [21] showed how to simultaneously bound the size of the bunches and clusters.

Lemma 4 ([21]). *Given an integer parameter $s > 0$, it is possible to construct a set A with expected size of $2s\log n$, such that $|C_A(w)| \leq 4n/s$, for every $w \in V$.*

We will also make use of the following well known Lemma.

Lemma 5 (Hitting set [4, 12]). *Let $1 \leq s \leq n$. Given k sets $S_1,\ldots,S_k$, where $S_i \subseteq V$ and $|S_i| \geq s$, for every $i \in \{1,\ldots,k\}$, we can find a set $H \subseteq V$ of size $\tilde{O}(n/s)$, such that $H \cap S_i \neq \emptyset$.*

Abraham *et al.* [2] introduced a nice coloring technique and used it to obtain a name-independent routing scheme with stretch 3 and routing tables of size $\tilde{O}(\sqrt{n})$. Recently, Abraham and Gavoille [1] used this coloring technique to obtain a labeled routing scheme with stretch $(2,1)$ and routing tables of size $\tilde{O}(n^{3/4})$. The next Lemma is implicit in [2, 1].

Lemma 6 (Coloring [2, 1]). *Let $0 < q \leq n$ be an integer. Let $S_1,\ldots,S_k$ be vertex sets each of size at least $\alpha q\log n$, where α is a large enough constant. There is a coloring function $c : V \to \{1,\ldots,q\}$ that satisfies the following:*

1. *For each $i \in \{1,\ldots,k\}$ and $j \in \{1,\ldots,q\}$, there exist $w \in S_i$ such that $c(w) = j$.*
2. *For each $j \in \{1,\ldots,q\}$ the number of vertices of color j is $O(n/q)$.*

Abraham *et al.* [2] showed that such a deterministic coloring function can be computed efficiently. For the sake of completeness we show that such a coloring function is formed, with high probability, from a random uniform coloring of the vertices of the graph. Assume that the number of sets k is polynomial in n. Each vertex is colored by one of the q colors with uniform probability. Let $S \in \{S_1,\ldots,S_k\}$ and $j \in \{1,\ldots,q\}$. Consider the event that there is no vertex in S with color j. The probability for this event to happen is bounded by $(1-1/q)^s$ which is at most $(1/n)^\alpha$. From the union bound we get that the first requirement holds with probability $1 - k\cdot q\cdot(1/n)^\alpha$. The expected number of vertices with the same color is n/q, and from Chernoff bound it follows that the second requirement holds with high probability. Using the union bound to bound the probability that one of the requirements does not hold we get that the above function satisfies, whp, both requirements.

Finally, to simplify the presentation we define $\tilde{x} = \alpha x\log n$, where $x > 0$ is an integer and α is a large enough constant of our choice.

3 New routing techniques

In this section we present two new techniques for routing between predefined vertex sets. We will use these techniques in the next sections to obtain our new routing schemes. The techniques are presented in two lemmas. We first show that for a vertex partition that satisfies certain properties it is possible to route efficiently between vertices of the same set in the partition.

Lemma 7. *Let $G = (V, E, w)$ be a weighted undirected graph, where $w : E \to \mathbb{R}$. Let $\mathcal{U} = \{U_1, \ldots, U_q\}$ be a partition of V into q sets, each of size $\tilde{O}(n/q)$. For every $\varepsilon > 0$, there is a routing scheme that uses a header of $O(\frac{1}{\varepsilon} \log n + \log^2 n / \log\log n)$-bit to route a message between any pair of vertices of the same set in partition $\mathcal{U}$, on a $(1+\varepsilon)$-stretch path. The routing table stored at each vertex is of size $\tilde{O}(\frac{1}{\varepsilon}(n/q) + q)$.*

Proof.
Preprocessing: For each $u \in V$, we store $B(u, \tilde{q})$. We compute a hitting set H as in Lemma 5 of size $\tilde{O}(n/q)$ that hits for each $u \in V$ the set $B(u, \tilde{q})$. For each vertex $w \in H$, we compute a shortest path tree $T(w)$ rooted at w that spans V. For each such $T(w)$, we store at each vertex $u \in V$ the tree routing information of $T(w)$ as needed for the tree routing schemes of Lemma 3. Since $|H| = \tilde{O}(n/q)$ and the tree routing information is of poly-logarithmic size, the total storage at each vertex for this purpose is $\tilde{O}(n/q)$.

Let $U \in \mathcal{U}$ and let $u \in U$. For every $v \in U$, we compute a sequence $\langle x_1, \ldots, x_{b'} \rangle$ of vertices. The sequence is computed using the following process:

Let $b = \lceil 2/\varepsilon \rceil$. Set $x_0 = u$ and $s = d(u, v)/b$. Let x_i $(i \geq 0)$ be the last vertex added to the sequence so far. If $v \in B(x_i, \tilde{q})$ then set x_{i+1} to v and stop. If $v \notin B(x_i, \tilde{q})$ then let (y_i, z_i) be an edge on a shortest path from x_i to v such that $y_i \in B(x_i, \tilde{q})$ and $z_i \notin B(x_i, \tilde{q})$. Proceed as follows:

- If $z_i = v$ then set $x_{i+1} = y_i$, $x_{i+2} = v$ and stop.
- If $d(x_i, z_i) < s$ then pick $w \in H$ such that $w \in B(x_i, \tilde{q})$, set $x_{i+1} = w$ and stop.
- Otherwise $(d(x_i, z_i) \geq s)$ set $x_{i+1} = y_i$, $x_{i+2} = z_i$ and continue to add vertices to the sequence with x_{i+2} being now the last added vertex.

In this process we add vertices to the sequence as long the progress towards v exceeds the threshold value s. Therefore, the process has at most $d(u, v)/s \leq b$ rounds. Since in each round at most 2 vertices are added to the sequence its total size is at most $2b$. At u we store for v the sequence $\langle x_1, \ldots, x_{b'} \rangle$. In case that $x_{b'} \neq v$ then it must be that $x_{b'} \in H$ and we also store at u the label of v in $T(x_{b'})$. Since $|U| = \tilde{O}(n/q)$, the total storage required at u for this information is $\tilde{O}(b \cdot (n/q))$.

Routing: Given a set $U \in \mathcal{U}$ and a pair of vertices $u, v \in U$ a message is routed from u to v as follows. First, u obtains the sequence $\langle x_1, \ldots, x_{b'} \rangle$ for v from its routing table and adds it to the message header[1]. It also sets x_1 to be a temporary target. Let u' be the current vertex on the routing path and let $x_i \neq u'$ be the temporary target. Assume that we are either in the case that $i < b'$ or in the case that $i = b'$ and $x_{b'} = v$. By the way $\langle x_1, \ldots, x_{b'} \rangle$ was constructed it follows that either $(u', x_i) \in E$ or $x_i \in B(u', \tilde{q})$. In the case that $(u', x_i) \in E$ then we can route from u' to x_i using a direct link[2]. In the case that $x_i \in B(u', \tilde{q})$, it follows from Lemma 2 that we can route to x_i on a shortest path. Each time a temporary target $x_i \neq v$ receives the message, it sets x_{i+1} to be the next temporary target. For the case that $i = b'$ and $x_{b'} \neq v$, it follows from the construction of $\langle x_1, \ldots, x_{b'} \rangle$ that $x_{b'} \in H$. The routing from $x_{b'-1}$ to v is done using the tree $T(x_{b'})$. This routing procedure requires a header of $O(\frac{1}{\varepsilon} \log n + \log^2 n / \log\log n)$-bit.

[1] In the case that $x_{b'} \neq v$, it adds also the label of v in $T(x_{b'})$ that was stored at u in the preprocessing.

[2] The standard routing scheme model of [18] assumes that at u given a neighbor v it is possible to obtain the link that connects u to v. We can even avoid this assumption by storing in the sequence edges instead of vertices when needed.

We now turn to analyze the stretch of the routing path. By the sequence construction, until $x_{b'-1}$, the message is routed on a shortest path. In the case that $x_{b'} = v$ then the message is routed from $x_{b'-1}$ to $x_{b'}$ on a shortest path as well. In the case that $x_{b'} \neq v$ the message is routed from $x_{b'-1}$ to v on $T(x_{b'})$, where $x_{b'} \in H \cap B(x_{b'-1}, \tilde{q})$. We have added $x_{b'}$ to the sequence because $d(x_{b'-1}, z_{b'-1}) < d(u,v)/b$. We also know that $z_{b'-1} \notin B(x_{b'-1}, \tilde{q})$, thus, $d(x_{b'-1}, x_{b'}) \leq d(x_{b'-1}, z_{b'-1}) < d(u,v)/b$. The length of the routing path is bounded by $d(u, x_{b'-1}) + d(x_{b'-1}, x_{b'}) + d(x_{b'}, v)$. Using the triangle inequality we bound $d(x_{b'}, v)$ with $d(x_{b'-1}, x_{b'}) + d(x_{b'-1}, v)$ and get:

$$d(u, x_{b'-1}) + d(x_{b'-1}, x_{b'}) + d(x_{b'}, v) \leq d(u, x_{b'-1}) + 2d(x_{b'-1}, x_{b'}) + d(x_{b'-1}, v).$$

Since $x_{b'-1}$ is on a shortest path between u and v it follows that $d(u, x_{b'-1}) + d(x_{b'-1}, v) = d(u,v)$ and we get:

$$d(u, x_{b'-1}) + 2d(x_{b'-1}, x_{b'}) + d(x_{b'-1}, v) \leq d(u,v) + 2d(u,v)/b.$$

Since $b = \lceil 2/\varepsilon \rceil$ we get that the message traverses a $(1+\varepsilon)$-stretch path. □

Next we extend Lemma 7 to the case that the possible destinations are from a subset W of V.

Lemma 8. *Let $G = (V, E, \omega)$ be a weighted undirected graph, where $\omega : E \to \mathbb{R}^+$. Let $\mathcal{W} = \{W_1, \ldots, W_q\}$ be a partition of $W \subseteq V$ into q sets, each of size $\tilde{O}(|W|/q)$. Let $\mathcal{U} = \{U_1, \ldots, U_q\}$ be a partition of V into q sets, such that $U \cap B(u, \tilde{q}) \neq \emptyset$, for every $U \in \mathcal{U}$ and $u \in V$. For every $\varepsilon > 0$, there is a routing scheme that uses a header size of $O(\frac{1}{\varepsilon}\log(nD))$-bit, and for every $1 \leq i \leq q$ routes a message from any vertex of U_i to any vertex of W_i on a $(1+\varepsilon)$-stretch path. The routing table stored at each vertex is of size $\tilde{O}(q + \frac{1}{\varepsilon} \cdot \log(D)(|W|/q))$.*

Proof.
Preprocessing:

We first assume that $\omega : E \to [1, M]$. Similarly to Lemma 7, for each vertex $u \in V$ we store $B(u, \tilde{q})$. It requires $\tilde{O}(q)$ space at each vertex.

Let $b = \lceil \frac{2}{\varepsilon} \rceil + 1$. For $1 \leq j \leq q$, let $U_j \in \mathcal{U}$ and let $W_j \in \mathcal{W}$. Given a vertex $u \in U_j$ and a vertex $w \in W_j$, we store at u a sequence of vertices that will be used to route towards w. Such a sequence has $O(\log(Mn))$ subsequences. A subsequence $\langle x_1, \ldots, x_{b'} \rangle$ is computed as follows.

Let x be a start vertex and let s be a threshold value that are given as an input to the computation. Set $x_0 = x$. Let x_i ($i \geq 0$) be the last vertex added to the subsequence so far. If $w \in B(x_i, \tilde{q})$ then set x_{i+1} to w and stop. If $w \notin B(x_i, \tilde{q})$ then let (y_i, z_i) be an edge on a shortest path from x_i to w such that $y_i \in B(x_i, \tilde{q})$ and $z_i \notin B(x_i, \tilde{q})$. Proceed as follows:

- If $z_i = w$ then set $x_{i+1} = y_i$, $x_{i+2} = w$ and stop.
- If $d(x_i, z_i) < s$ then there is a vertex $z \in B(x_i, \tilde{q}) \cap U_j$, set $x_{i+1} = z$ and stop.
- Otherwise ($d(x_i, z_i) \geq s$) set $x_{i+1} = y_i$, $x_{i+2} = z_i$. If the subsequence is of size $2b$ stop, else continue to add vertices to the subsequence with x_{i+2} being now the last added vertex.

Now, the sequence stored at u for w is computed as follows. Let $P(u,w) = \{u = u_0, u_1, u_2, \ldots, u_t = w\}$ be a shortest path between u and w. We start by adding u_1 to the sequence. If $u_1 = w$ we stop, otherwise, we add u_2 to the sequence as well. Next, if $u_2 \neq w$, we start to produce subsequences and add them to the sequence. The first subsequence is produced using $x = u_2$ as the start vertex and $s = 2/b$ as the threshold value. Let s' be the threshold value that was used to produce the last subsequence added so far to the sequence and let w' be its last vertex. If this subsequence has exactly $2b$ vertices and $w' \neq w$ then a new

subsequence is produced and added to the sequence using $x = w'$ as the start vertex and $s = 2s'$ as the threshold value.

Each subsequence added to the sequence has at most $2b$ vertices. In every subsequence, beside maybe the last one, all the vertices are on $P(u, w)$. Moreover, if the ith subsequence added to the sequence is not the last subsequence its last vertex is at a distance of at least 2^i from its first vertex. Since the length of every shortest path is less than Mn the total number of subsequences cannot exceed $\log(nM)$ and therefore a sequence can have at most $2b\log(Mn)+2$ vertices. Each set in $\mathcal{W}$ is of size $|W|/q$, thus, each $u \in U_j$ stores $|W|/q$ sequences which results in $O(b(\log Mn)|W|/q) = \tilde{O}(\frac{1}{\varepsilon}(\log M)|W|/q)$ in total.

Next, we show that the preprocessing algorithm works also on graphs with non-negative real edge weights. Let $G = (V, E, \omega)$ be a weighted graph with $\omega : E \to \mathbb{R}^+$ and let D be its normalized diameter. Let $E' = \{(u, v) \in E \mid \omega(u, v) = d(u, v)\}$. Let $\omega_{\max} = max\{\omega(u, v) \mid (u, v) \in E'\}$, and let $\omega_{\min} = min\{\omega(u, v) \mid (u, v) \in E'\}$. For every $(u, v) \in E'$ let $\omega'(u, v) = \omega(u, v)/\omega_{\min}$. It is easy to see that $\omega' : E' \to [1, M]$, where $M = \frac{\omega_{\max}}{\omega_{\min}}$. Notice that all shortest paths in $G' = (V, E', w')$ are also shortest paths in G. Applying the preprocessing algorithm on G' we get that each $u \in U_j$ stores $|W|/q$ sequences each of size $O(b\log(Mn))$. Since $\omega_{\max} \leq \max_{u,v} d(u, v)$ and $\omega_{\min} = \max_{u \neq v} d(u, v)$ it follows that $M = \frac{\omega_{\max}}{\omega_{\min}} \leq D$. Therefore, the total space stored at a vertex is $\tilde{O}(\frac{1}{\varepsilon}(\log D)|W|/q)$.

Routing: Given a set $U_j \in \mathcal{U}$ and a set $W_j \in \mathcal{W}$, and a pair of vertices $u \in U_j$ and $w \in W_j$, we show how to route a message from u to w. We route towards the last vertex of the sequence stored at u for w in the same manner as in Lemma 7. It follows from the sequence construction that all its vertices, beside maybe the last, are on a shortest path between u and w. The last vertex is either w or a vertex that belongs to U_j. Let $r_0 = u$. For $i \geq 0$, let r_{i+1} be the last vertex of the sequence stored at r_i for w, and let r'_i be the vertex that precedes r_{i+1} in the sequence. If $r_{i+1} \neq w$ then $r_{i+1} \in U_j$. Let $\alpha_i = d(r_i, r'_i)$.

Claim 9. *For every $i \geq 0$, $d(r_{i+1}, w) < d(r_i, w)$. Moreover, $d(r_{i+1}, w) \leq d(r_i, w) - (\alpha_i - \alpha_i/b)$.*

Proof. Let $i \geq 0$. Consider the routing of a message from r_i to r_{i+1}. If $r_{i+1} = w$ then the message is routed on a shortest path, therefore $d(r_{i+1}, w) = d(w, w) < d(r_i, w)$ and also $d(r_{i+1}, w) = 0 \leq d(r_i, w) - d(r_i, r'_i) \leq d(r_i, w) - (\alpha_i - \alpha_i/b)$. Otherwise, we route from $r_i \in U_j$ to a vertex $r_{i+1} \in U_j$. The sequence stored at $r_i \in U_j$ for $w \in W_j$ contains the first two vertices of a shortest path from r_i to w and then at least one subsequence. Let $k \geq 1$ be the number of subsequences of this sequence. The last vertex of this sequence is r_{i+1}, and r'_i is the vertex that precedes r_{i+1} in the sequence. From the sequence construction it follows that r'_i is on a shortest path between r_i and w. Moreover, r'_i is either the last vertex of the $(k-1)$th subsequence or a vertex of the kth subsequence or the second vertex of the sequence.

The message is routed from r_i to r'_i, and then from r'_i to r_{i+1}. From the way that subsequences are produced it follows that $\alpha_i = d(r_i, r'_i) \geq 2 + (2 + 4 + \ldots + 2^{k-1}) = 1 + (1 + 2 + \ldots + 2^{k-1}) = 2^k$. Moreover, it follows from the threshold used to produce the kth subsequence that $d(r'_i, r_{i+1}) \leq 2^k/b \leq \alpha_i/b$. Since r'_i is on a shortest path between r_i and w it follows that $d(r'_i, w) = d(r_i, w) - d(r_i, r'_i)$. From the triangle inequality it follows that $d(r_{i+1}, w) \leq d(r_{i+1}, r'_i) + d(r'_i, w)$. We get:

$$d(r_{i+1}, w) \leq d(r_{i+1}, r'_i) + d(r'_i, w) = d(r_{i+1}, r'_i) + d(r_i, w) - d(r_i, r'_i) \leq d(r_i, w) - (\alpha_i - \alpha_i/b),$$

as required.

□

Next, we show that the stretch of the routing path is $(1 + \varepsilon)$. From Claim 9 it follows that the message eventually reaches a vertex r_t that w is the last vertex in its sequence. For every $0 \leq i < t$, the message is routed from r_i to r_{i+1} on a path of length at most $\alpha_i + \alpha_i/b$. For $i = t$ the message is routed on

a shortest path between r_t and w. Let $A = \sum_{i=0}^{t-1} \alpha_i$. The message traverses a path of length at most $A(1+1/b) + d(r_t, w)$.

From Claim 9 it follows that $d(r_t, w) \leq d(r_{t-1}, w) - (\alpha_{t-1} - \alpha_{t-1}/b)$. By applying Claim 9 recursively we get $d(r_t, w) \leq d(u, w) - \sum_{i=0}^{t-1}(\alpha_i - \alpha_i/b) = d(u, w) - A(1-1/b)$. From this we get that $A(1-1/b) \leq d(u, w) - d(r_t, w) \leq d(u, w)$ and $A \leq \frac{b}{b-1} d(u, w)$. We can now bound the length of the path that the message traverses:

$$\begin{aligned} A(1+1/b) + d(r_t, w) &\leq A(1+1/b) + d(u, w) - A(1-1/b) \\ &\leq d(u, w) + 2A/b \\ &\leq d(u, w) + 2d(u, w)/(b-1) \\ &= (1 + \frac{2}{b-1}) d(u, w) \\ &\leq (1+\varepsilon) d(u, w) \end{aligned}$$

□

4 Applications for small stretch routing schemes

In this section we show how to obtain small stretch routing schemes, that almost match the corresponding state-of-the-art distance oracles, using the routing techniques from the previous section. To exemplify the strength of our techniques we first show as a warm-up a new $(3+\varepsilon)$-stretch labeled routing scheme with $\tilde{O}(\frac{1}{\varepsilon})$-bit header size and $\tilde{O}(\frac{1}{\varepsilon}\sqrt{n})$-space routing tables. Using the hash function presented in [2] it is easy to modify this routing scheme to be name-independent. Then we turn to present the main results of this section, our $(2+\varepsilon, 1)$-stretch and $(5+\varepsilon)$-stretch routing schemes.

Stretch $3+\varepsilon$. Let $q = \sqrt{n}$. Let c be a coloring function with q colors obtained using Lemma 6 with respect to the sets $B(u, \tilde{q})$, for every $u \in V$. Let $\mathcal{U} = \{U_1, \ldots, U_q\}$ be the color sets induced by c. It follows that $\mathcal{U}$ is a partition and every $U \in \mathcal{U}$ has $\tilde{O}(\sqrt{n})$ vertices. We use Lemma 7 with $\mathcal{U}$. The information stored for this at every vertex is of size $\tilde{O}(\frac{1}{\varepsilon}\sqrt{n})$. Every $u \in V$ stores $B(u, \tilde{q})$. It follows from Lemma 6 that $B(u, \tilde{q})$ contains a vertex of each color. For every $i \in \{1, \ldots, q\}$ we store at u a vertex from $B(u, \tilde{q})$ of color i. For every $v \in V$ the label of v contains v and $c(v)$. A message is routed from u to v as follows. If $v \in B(u, \tilde{q})$ it follows from Lemma 2 that we can route the message from u to v on a shortest path. If $v \notin B(u, \tilde{q})$ then, using the information that is saved at u, we route to $w \in B(u, \tilde{q})$ with color $c(v)$. Again from Lemma 2 it follows that we can route from u to w on a shortest path. We route from w to v using Lemma 7. We now bound the stretch. From Lemma 7 it follows that the message is sent from w to v on a path of length at most $(1+\varepsilon)d(w, v)$. The message traverses a path of length at most $d(u, w) + (1+\varepsilon)d(w, v)$. From the triangle inequality it follows that $d(w, v) \leq d(w, u) + d(u, v)$. Also, since $v \notin B(u, \tilde{q})$ and $w \in B(u, \tilde{q})$ it follows that $d(u, w) \leq d(u, v)$. Therefore, $d(w, v) \leq 2d(u, v)$, and we get that $d(u, w) + (1+\varepsilon)d(w, v) \leq (3+2\varepsilon)d(u, v)$.

Stretch $(2+\varepsilon, 1)$. We now present a $(2+\varepsilon, 1)$-stretch routing scheme with $o(\log^2 n)$-bit labels, $\tilde{O}(1/\varepsilon)$-bit headers and routing tables of $\tilde{O}(\frac{1}{\varepsilon} n^{2/3})$ size. This almost matches the $(2, 1)$-stretch distance oracle with $\tilde{O}(n^{5/3})$ total space of [19].

Theorem 10. *There is a routing scheme for unweighted undirected graphs with $o(\log^2 n)$-bit labels, that uses $\tilde{O}(\frac{1}{\varepsilon} n^{2/3})$ space at each vertex and $\tilde{O}(1/\varepsilon)$-bit header size, to route a message between any pair of vertices on a $(2+\varepsilon, 1)$-stretch path.*

Proof. Let $q = n^{1/3}$. Every $u \in V$ stores $B(u, \tilde{q})$. We compute using Lemma 4 a set $A \subset V$ of size $\tilde{O}(n^{2/3})$, such that $|C_A(w)| = O(n^{1/3})$ for every $w \in V$. For each $u \in C_A(w)$ we store the routing information for a tree routing scheme of $T_{C_A}(w)$ at u. This information is of poly-logarithmic size. Every vertex u is contained in at most $O(n^{1/3})$ cluster trees as $|B_A(u)| = O(n^{1/3})$. Therefore, the storage required at each vertex u for this is $\tilde{O}(n^{1/3})$.

Let $w \in V$. For each $v \in C_A(w)$ we store at w the label of v in the tree routing scheme of $T_{C_A}(w)$. Since $|C_A(w)| = O(n^{1/3})$ the storage required at w for these labels is $\tilde{O}(n^{1/3})$.

For each $w \in A$ let $T(w)$ be a shortest path tree rooted at w that spans V. For every $v \in V$ we store routing information of a tree routing scheme of $T(w)$. This allows to route from a vertex u to a vertex v on $T(w)$. Since $|A| = \tilde{O}(n^{2/3})$, each vertex stores $\tilde{O}(n^{2/3})$ information for this purpose.

Let $u \in V$. For every $v \in V$, if $B(u, \tilde{q}) \cap B_A(v) \neq \emptyset$ we store in a hash table at u in the entry of v a vertex $w \in \arg\min_{w'}\{d(u, w') + d(w', v) \mid w' \in B(u, \tilde{q}) \cap B_A(v)\}$. There are $\tilde{O}(n^{1/3})$ vertices in $B(u, \tilde{q})$ each with a cluster of $O(n^{1/3})$ vertices. Let $w \in B(u, \tilde{q})$. For every $v \in C_A(w)$ it holds that $B(u, \tilde{q}) \cap B_A(v) \neq \emptyset$ therefore the storage required for this at every vertex is $\tilde{O}(n^{2/3})$.

Let c be a coloring function with q colors obtained using Lemma 6 with respect to the sets $B(u, \tilde{q})$, for every $u \in V$. Let $\mathcal{U} = \{U_1, \ldots, U_q\}$ be the color sets induced by c. It follows that $\mathcal{U}$ is a partition and every $U \in \mathcal{U}$ has $\tilde{O}(n/q) = \tilde{O}(n^{2/3})$ vertices. We use Lemma 7 with the partition $\mathcal{U}$. The information stored at the routing table of every vertex for this is $\tilde{O}(\frac{1}{\varepsilon}n^{2/3})$.

It follows from Lemma 6 that $B(u, \tilde{q})$ contains a vertex of each color. For every $i \in \{1, \ldots, q\}$ we store at u a vertex from $B(u, \tilde{q})$ of color i and its distance from u.

Finally, the label of $v \in V$ contains the following information: v, $c(v)$, $p_A(v)$, $d(v, p_A(v))$ and the label of v in $T(p_A(v))$.

Routing: A message is routed from u to v as follows. At u we check in constant time if $B(u, \tilde{q}) \cap B_A(v) \neq \emptyset$. If this is the case we obtain in constant time the vertex $w \in \arg\min_{w'}\{d(u, w') + d(w', v) \mid w' \in B(u, \tilde{q}) \cap B_A(v)\}$ that was saved at u. From Lemma 2 it follows that we can route the message from u to w on a shortest path. From w to v we can use the tree routing scheme of $T_{C_A}(w)$ as the label of v in the this tree routing scheme is stored at w. Thus, the message is routed on a path of length $d(u, w) + d(w, v)$. Next, we show that $d(u, v) = d(u, w) + d(w, v)$. Assume, towards a contradiction, that w is not on a shortest path between u and v and let P be a shortest path between u and v. There is a vertex $u' \in B(u, \tilde{q}) \cap P$ such that $d(u, u') = r_u$, and there is a vertex $v' \in B_A(v) \cap P$ such that $d(v, v') = d(v, p_A(v)) - 1$ (notice that $B(u, \tilde{q}) \cap B_A(v) \neq \emptyset$ implies that $v \notin A$ and $d(v, p_A(v)) - 1 \geq 0$). Since the intersection is not on a vertex from P it must be that $d(u, v) \geq r_u + d(v, v') + 1$. On the other hand $d(u, w) \leq r_u + 1$ and $d(v, w) \leq d(v, v')$ and we get that $d(u, w) + d(v, w) \leq d(u, v)$, which contradicts the assumption that w is not on a shortest path between u and v.

Consider the case that $B(u, \tilde{q}) \cap B_A(v) = \emptyset$. We obtain $c(v)$, $p_A(v)$ and $d_A(p_A(v))$ from the label of v. A vertex $w \in B(u, \tilde{q})$ with $c(w) = c(v)$ is stored at u with $d(u, w)$. If $d(v, p_A(v)) \leq d(u, w)$ then the message is routed from u to v on $T(p_A(v))$. If $d(v, p_A(v)) > d(u, w)$, we route on a shortest path from u to w using Lemma 2 and then from w to v on a path of length at most $(1 + \varepsilon)d(w, v)$ using Lemma 7.

If $v \in A$, it follows that $d(v, p_A(v)) = 0 \leq d(u, w)$ and the message is routed on a shortest path to v in $T(v) = T(p_A(v))$. Assume $v \notin A$. As $B(u, \tilde{q}) \cap B_A(v) = \emptyset$ it follows that $d(u, v) \geq r_u + (d(v, p_A(v)) - 1) + 1 = r_u + d(v, p_A(v))$. If $d(v, p_A(v)) \leq d(u, w)$ then the message is sent on a path of length $d(u, p_A(v)) + d(p_A(v), v)$. From the triangle inequality it follows that $d(u, p_A(v)) \leq d(u, v) + d(p_A(v), v)$ and thus $d(u, p_A(v)) + d(p_A(v), v) \leq d(u, v) + 2d(p_A(v), v)$. Since $d(u, w) \leq r_u + 1$ we get that $d(v, p_A(v)) - 1 \leq r_u$. Combining this with the fact that $d(u, v) \geq r_u + d(v, p_A(v))$ we get that $d(u, v) + 1 \geq 2d(v, p_A(v))$. Therefore, $d(u, v) + 2d(p_A(v), v)$ is at most $2d(u, v) + 1$. If $d(u, w) < d(v, p_A(v))$ the message is sent on a path of length at most $d(u, w) + (1 + \varepsilon)d(w, v)$. From the triangle inequality it follows that $d(w, v) \leq d(u, w) + d(u, v)$

and thus $d(u,w)+(1+\varepsilon)d(w,v) \leq (1+\varepsilon)d(u,v)+(2+\varepsilon)d(u,w)$. We have $d(u,w) \leq d(v,p_A(v))-1$. We also have $d(u,w) \leq r_u+1$. Combining the last two inequalities with $d(u,v) \geq r_u+d(v,p_A(v))$ we get that $2d(u,w) \leq d(u,v)$. Therefore, the message is routed on a path of length at most $(2+2\varepsilon)d(u,v)$. □

Stretch $5+\varepsilon$. We now present a $(5+\varepsilon)$-stretch routing scheme that uses labels of size $(4\log n)$-bit, $\tilde{O}(1/\varepsilon \log D)$-bit headers and routing tables of $\tilde{O}(\frac{1}{\varepsilon}n^{1/3}\log D)$ size. This almost matches the 5-stretch distance oracle with $\tilde{O}(n^{4/3})$ total space of [22]. The current best routing scheme with routing tables of $\tilde{O}(n^{1/3})$ size has a stretch of 7 and $o(\log^2 n)$-bit labels. This routing scheme is obtained by setting $k=3$ in the $(4k-5)$-stretch routing scheme of Thorup and Zwick [21].

Theorem 11. *Let $G=(V,E,\omega)$ be a weighted undirected graph, where $\omega : E \to \mathbb{R}^+$. There is a routing scheme for G with $O(\log n)$-bit labels that uses $\tilde{O}(\frac{1}{\varepsilon}n^{1/3}\log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon}\log D)$-bit headers, to route a message between any pair of vertices on a $(5+\varepsilon)$-stretch path.*

Proof.
Preprocessing: Let $q=n^{1/3}$. Every $u \in V$ stores $B(u,\tilde{q})$. We compute using Lemma 4 a set $A \subset V$ of size $\tilde{O}(n^{2/3})$, such that $|C_A(w)| = O(n^{1/3})$ for every $w \in V$. Let $w \in V$, for each $v \in C_A(w)$ we store at w the label of v in the tree routing scheme of $T_{C_A}(w)$. Since $|C_A(w)| = O(n^{1/3})$ the storage required at w is $\tilde{O}(n^{1/3})$. For each $u \in C_A(w)$ we store at u the routing information for a tree routing scheme of $T_{C_A}(w)$. This information is of poly-logarithmic size. Notice that $|B_A(u)| = O(n^{1/3})$, thus u is contained in at most $O(n^{1/3})$ trees and the storage required at u for this is $\tilde{O}(n^{1/3})$.

Let c be a coloring function with q colors obtained using Lemma 6 with respect to the sets $B(u,\tilde{q})$, for every $u \in V$. Let $\mathcal{U}=\{U_1,\ldots,U_q\}$ be the color sets induced by c. Let $\mathcal{W}=\{W_1,\ldots,W_q\}$ be an arbitrary partition of the set A into q sets each with at most $|A|/q$ vertices. For $w \in A$, let $\alpha(w)$ be the index of the set in partition $\mathcal{W}$ that contains w, that is $w \in W_{\alpha(w)}$. We use Lemma 8 with partitions $\mathcal{U}$ and $\mathcal{W}$. The required storage at each vertex is $\tilde{O}(q+\frac{1}{\varepsilon}(\log D)|A|/q) = \tilde{O}(\frac{1}{\varepsilon}(\log D)n^{1/3})$.

It follows from Lemma 6 that $B(u,\tilde{q})$ contains a vertex of each color. For every $i \in \{1,\ldots,q\}$ we store at u a vertex from $B(u,\tilde{q})$ of color i. The storage required for this at u is $\tilde{O}(n^{1/3})$.

Let $v \in V$ and let $(p_A(v),z)$ be the first edge on a shortest path from $p_A(v)$ to v. Notice that $v \in C_A(z)$. The label of v contains the following information: v, $p_A(v)$, $\alpha(p_A(v))$ and $(p_A(v),z)$.

Routing: A message is routed from u to v as follows. At u we check if $v \in B(u,\tilde{q})$. If this is the case then from Lemma 2 it follows that we can route the message from u to v on a shortest path. If not we check if $v \in C_A(u)$. If this is the case then the label of v in the tree routing scheme of $T_{C_A}(u)$ is stored at u. The message is again routed on a shortest path between u and v.

In the case that $v \notin B(u,\tilde{q})$ and $v \notin C_A(u)$ the message is routed on a shortest path to $w \in B(u,\tilde{q})$ that satisfies $c(w)=\alpha(p_A(v))$. The message is routed from w to $p_A(v)$ using Lemma 8. At $p_A(v)$ we obtain $(p_A(v),z)$ from the label of v and forward the message to z. At z we have the label of v in the tree routing scheme of $T_{C_A}(z)$ since $v \in C_A(z)$. Using this label and the tree routing scheme of $T_{C_A}(z)$ the message is routed from z to v. The message is routed on a path of length at most $d(u,w)+(1+\varepsilon)d(w,p_A(v))+d(p_A(v),v)$. Since $v \notin B(u,\tilde{q})$ it follows that $d(u,w) \leq d(u,v)$. Since $v \notin C_A(u)$ it follows that $d(v,p_A(v)) \leq d(u,v)$. From the triangle inequality it follows that $d(w,p_A(v)) \leq d(u,w)+d(u,v)+d(p_A(v),v) \leq 3d(u,v)$. We get $d(u,w)+(1+\varepsilon)d(w,p_A(v))+d(p_A(v),v) \leq (5+3\varepsilon)d(u,v)$. □

5 Generalized routing schemes

Stretch $(3\pm 2/\ell+\varepsilon,2)$. We now present a $(3\pm 2/\ell+\varepsilon,2)$-stretch routing scheme that uses $O(\ell\log n)$-bit labels, $\tilde{O}(1/\varepsilon)$-bit headers and routing tables of $\tilde{O}(\ell\frac{1}{\varepsilon}n^{\ell/(2\ell\pm 1)})$ size. This almost matches the distance oracles presented in [20] for weighted graphs with $(3\pm 2/\ell)$-stretch and $\tilde{O}(\ell m^{1+\ell/(2\ell\pm 1)})$ total space.

We start with a simple Lemma given in [20] and repeated here for completeness.

Lemma 12. *Let $\ell>0$ be an integer. Let $\{x_i\}_{i=0}^{\ell}$ and $\{y_i\}_{i=0}^{\ell}$ be series of real numbers from $[0,1]$, such that $x_0=y_0=0$ and $x_i+y_{\ell-i}\leq 1$, for every $i\in\{0,\ldots,\ell\}$. There exists $i\in\{0,\ldots,\ell-1\}$ that satisfies $x_i+y_{\ell-i-1}\leq 1-1/\ell$.*

We now turn to present the $(3-2/\ell+\varepsilon,2)$-stretch routing scheme.

Theorem 13. *Let $\ell>1$ be an integer. There is a routing scheme for unweighted undirected graphs that uses $O(\ell\log n)$-bit labels, $\tilde{O}(\ell\frac{1}{\varepsilon}n^{\ell/(2\ell-1)})$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon})$ header size, to route a message between any pair of vertices on a $(3-2/\ell+\varepsilon,2)$-stretch path.*

Proof.
Preprocessing: Let $q=n^{1/(2\ell-1)}$ and let $i\in\{0,\ldots,\ell\}$. For every $u\in V$, let $B_i(u)=B(u,\tilde{q}^i)$. Every $u\in V$ stores $B_\ell(u)$.

For every $i\in\{0,\ldots,\ell\}$:

- We compute using Lemma 4 a set $L_i\subseteq V$ of size $\tilde{O}(q^{2\ell-i-1})$, such that $|C_{L_i}(w)|=O(n/q^{2\ell-i-1})=O(q^i)$, for every $w\in V$. For each $u\in C_{L_i}(w)$ we store at u the routing information for a tree routing scheme of $T_{C_{L_i}}(w)$. This information is of poly-logarithmic size. Notice that $|B_{L_i}(u)|=O(q^i)$, thus u is contained in at most $O(q^i)$ trees and the storage required at u for this is $\tilde{O}(q^i)$.
- Let $w\in V$. For each $v\in C_{L_i}(w)$ we store at w the label of v in the tree routing scheme of $T_{C_{L_i}}(w)$. Since $|C_{L_i}(w)|=O(q^i)$ the storage required at w is $\tilde{O}(q^i)$.
- Let $u\in V$. For every $v\in V$, if $B_i(u)\cap B_{L_{\ell-i}}(v)\neq\emptyset$ we store in a hash table i at u in the entry of v a vertex $w\in\arg\min_{w'\in B_i(u)\cap B_{L_{\ell-i}}(v)}\{d(u,w')+d(w',v)\}$. Since $|C_{L_{\ell-i}}(w)|=O(q^{\ell-i})$ and $|B_i(u)|=\tilde{O}(q^i)$, the storage required for this at every vertex is $\tilde{O}(q^\ell)$.

For every $i\in\{0,\ldots,\ell-1\}$:

- Let c_i be a coloring function obtained using Lemma 6 with respect to the sets $B_i(u)$ for every $u\in V$, and let $\mathcal{U}^i=\{U_1^i,\ldots,U_{q^i}^i\}$ be a partition of V into q^i color sets induced by c_i. Let $j=\ell-i-1$ and let $\mathcal{W}^j=\{W_1^j,\ldots,W_{q^i}^j\}$ be an arbitrary partition of L_j into q^i sets each with at most $|L_j|/q^i=q^{2\ell-(\ell-i-1)-1}/q^i=q^{\ell+i}/q^i=q^\ell$ vertices. For $w\in L_j$, let $\alpha_j(w)$ be the index of the set in $\mathcal{W}^j$ that contains w, that is $w\in W^j_{\alpha_j(w)}$. We use Lemma 8 with the partitions $\mathcal{U}^i$ and $\mathcal{W}^j$. The required storage at each vertex is $\tilde{O}(\frac{1}{\varepsilon}|L_j|/q^i+q^i)=\tilde{O}(\frac{1}{\varepsilon}q^\ell)$.
- At u we store for each $j\in\{1,\ldots,q^i\}$ a vertex $w\in B_i(u)$ with $c_i(w)=j$. By Lemma 6 such a vertex must exist. The storage required for this at u is $\tilde{O}(q^i)$.

In the label of $v\in V$ we store v and $\{p_{L_i}(v),\alpha_i(p_{L_i}(v)),d(v,p_{L_i}(v)),(p_{L_i}(v),v_i')\}_{i=0}^{\ell-1}$, where $(p_{L_i}(v),v_i')$ is the first edge on a shortest path from $p_{L_i}(v)$ to v.

Routing: A message is routed from u to v as follows. At u we check if there exists $i \in \{0,\ldots,\ell\}$ such that $B_i(u) \cap B_{L_{\ell-i}}(v) \neq \emptyset$. If this is the case we can obtain the vertex $w \in B_i(u) \cap B_{L_{\ell-i}}(v)$ that was saved at u. From Lemma 2 it follows that we can route the message from u to w on a shortest path. From w to v we can use the tree routing scheme of $T_{C_{L_{\ell-i}}}(w)$ as the label of v in the this tree routing scheme is stored at w. Thus, the message is routed on a path of length $d(u,w)+d(w,v)$. The vertex w is on a shortest path between u and v. The proof for that is identical to the proof given in Theorem 10 for the case of intersection.

Consider the case that for every $i \in \{0,\ldots,\ell\}$ we have $B_i(u) \cap B_{L_{\ell-i}}(v) = \emptyset$. Let $a_i = r_u(\tilde{q}^i)$ and let $b_i = d(v, p_{L_i}(v)) - 1$ if $v \notin L_i$ and $b_i = 0$, otherwise. Let $j \in \arg\min_{i\in\{0,\ldots,\ell-1\}}\{a_i + b_{\ell-i-1}\}$. If there is more than one index that achieves the minimum we take the index of the highest value. Let $k = \ell - j - 1$. We obtain $\alpha_k(p_{L_k}(v))$ from the label of v and route on a shortest path to the vertex $w \in B_j(u)$ with $c_j(w) = \alpha_k(p_{L_k}(v))$, that is stored at u. From w the message is routed to $p_{L_k}(v)$ using Lemma 8. From $p_{L_k}(v)$ the message is forwarded to v'_k and then using the tree routing scheme of $T_{C_{L_k}}(v'_k)$ it is routed to v. The total length of the path is at most $d(u,w) + (1+\varepsilon)d(w, p_{L_k}(v)) + d(p_{L_k}(v), v)$. Let $d(u,v) = \Delta$.

We first consider the degenerated case of $\Delta = 1$. Since there is no intersection we have $B_\ell(u) \cap B_{L_0}(v) = \emptyset$, where $B_{L_0}(v) = \{v\}$. Therefore, $a_{\ell-1} \leq a_\ell = 0$ and $b_0 = 0$, and by the selection rule of j it must be that $j = \ell - 1$. The message is routed to $w \in B_{\ell-1}(u)$ and from there to $p_{L_0}(v) = v$. The length of this path is at most $3 + \varepsilon$.

Consider now the case that $\Delta > 1$. From the triangle inequality it follows that $d(w, p_{L_k}(v)) \leq d(u,w) + d(u,v) + d(v, p_{L_k}(v))$. If $j = 0$ then $u = w$, $a_0 = 0$, and we need to bound $(2+\varepsilon)(b_{\ell-1}+1) + (1+\varepsilon)\Delta$. If $j > 0$ then $d(u,w) \leq a_j + 1$, and we need to bound $(2+\varepsilon)(a_j + b_k + 2) + (1+\varepsilon)\Delta$.

For every $i \in \{0,\ldots,\ell\}$ we have $B_i(u) \cap B_{L_{\ell-i}}(v) = \emptyset$ thus $a_i + b_{\ell-i} \leq \Delta - 1$. Let $x_0 = a_0/\Delta$ and $x_i = (a_i+1)/\Delta$, for $i > 0$. Let $y_i = b_i/\Delta$. Since $x_0 = y_0 = 0$ and $x_i + y_{\ell-i} \leq 1$ we can apply Lemma 12 with the series $\{x_i\}_{i=0}^{\ell}$ and $\{y_i\}_{i=0}^{\ell}$.

For $j = 0$, we get from Lemma 12 that $a_0 + b_{\ell-1} = b_{\ell-1} \leq (1 - 1/\ell)\Delta$. Therefore,

$$\begin{aligned}(2+\varepsilon)(b_{\ell-1}+1) + (1+\varepsilon)\Delta &\leq (2+\varepsilon)((1-1/\ell)\Delta + 1) + (1+\varepsilon)\Delta \\ &= \Delta((1+\varepsilon) + (1-1/\ell)(2+\varepsilon)) + 2 + \varepsilon \\ &= \Delta(3 + 2\varepsilon - (2+\varepsilon)/\ell) + 2 + \varepsilon \\ &\leq \Delta(3 + 3\varepsilon - (2+\varepsilon)/\ell) + 2.\end{aligned}$$

If $j > 0$, it follows from Lemma 12 that $a_j + 1 + b_{\ell-j-1} = a_j + 1 + b_k \leq (1 - 1/\ell)\Delta$. Therefore,

$$\begin{aligned}(2+\varepsilon)(a_j + b_k + 2) + (1+\varepsilon)\Delta &\leq (2+\varepsilon)((1-1/\ell)\Delta + 1) + (1+\varepsilon)\Delta \\ &\leq \Delta(3 + 3\varepsilon - (2+\varepsilon)/\ell) + 2.\end{aligned}$$

□

The routing scheme of $(3 + 2/\ell + \varepsilon, 2)$-stretch is very similar to the one presented above. For the sake of completeness we provide it here with all the details.

As before we repeat a Lemma that was given in [20].

Lemma 14. *Let $\ell > 0$ be an integer. Let $\{x_i\}_{i=0}^{\ell}$ and $\{y_i\}_{i=0}^{\ell}$ be series of real numbers from $[0,1]$, such that $x_0 = y_0 = 0$ and $x_i + y_{\ell-i} \leq 1$, for every $i \in \{0,\ldots,\ell\}$. There exists $i \in \{0,\ldots,\ell-1\}$ that satisfies $x_{i+1} + y_{\ell-i} \leq 1 + 1/\ell$.*

We now turn to present the $(3+2/\ell+\varepsilon, 2)$-stretch routing scheme.

Theorem 15. *Let $\ell > 1$ be an integer. There is a routing scheme for unweighted undirected graphs that uses $O(\ell \log n)$-bit labels, $\tilde{O}(\ell \frac{1}{\varepsilon} n^{\ell/(2\ell+1)})$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon})$ headers, to route a message between any pair of vertices on a $(3+2/\ell+\varepsilon, 2)$-stretch path.*

Proof.
Preprocessing: Let $q = n^{1/(2\ell+1)}$ and let $i \in \{0,\ldots,\ell\}$. For every $u \in V$, let $B_i(u) = B(u, \tilde{q}^i)$. Every $u \in V$ stores $B_\ell(u)$. For every $i \in \{0,\ldots,\ell\}$:

- We compute using Lemma 4 a set $L_i \subseteq V$ of size $\tilde{O}(q^{2\ell-i+1})$, such that $|C_{L_i}(w)| = O(n/q^{2\ell-i+1}) = O(q^i)$, for every $w \in V$. For each $u \in C_{L_i}(w)$ we store at u the routing information for a tree routing scheme of $T_{C_{L_i}}(w)$. This information is of poly-logarithmic size. Notice that $|B_{L_i}(u)| = O(q^i)$, thus u is contained in at most $O(q^i)$ trees and the storage required at u for this is $\tilde{O}(q^i)$.
- Let $w \in V$. For each $v \in C_{L_i}(w)$ we store at w the label of v in the tree routing scheme of $T_{C_{L_i}}(w)$. Since $|C_{L_i}(w)| = O(q^i)$ the storage required at w is $\tilde{O}(q^i)$.
- Let $u \in V$. For every $v \in V$, if $B_i(u) \cap B_{L_{\ell-i}}(v) \neq \emptyset$ we store in a hash table i at u in the entry of v a vertex $w \in \arg\min_{w' \in B_i(u) \cap B_{L_{\ell-i}}(v)}\{d(u,w') + d(w',v)\}$. Since $|C_{L_{\ell-i}}(w)| = O(q^{\ell-i})$ and $|B_i(u)| = \tilde{O}(q^i)$, the storage required for this at every vertex is $\tilde{O}(q^\ell)$.

For every $i \in \{1,\ldots,\ell\}$:

- Let c_i be a coloring function obtained using Lemma 6 with respect to the sets $B_i(u)$ for every $u \in V$, and let $\mathcal{U}^i = \{U_1^i,\ldots,U_{q^i}^i\}$ be a partition of V into q^i color sets induced by c_i. Let $j = \ell - i + 1$ and let $\mathcal{W}^j = \{W_1^j,\ldots,W_{q^i}^j\}$ be an arbitrary partition of L_j into q^i sets each with at most $|L_j|/q^i = q^{2\ell-(\ell-i+1)+1}/q^i = q^{\ell+i}/q^i = q^\ell$ vertices. For $w \in L_j$, let $\alpha_j(w)$ be the index of the set in $\mathcal{W}^j$ that contains w, that is $w \in W_{\alpha_j(w)}^j$. We use Lemma 8 with the partitions $\mathcal{U}^i$ and $\mathcal{W}^j$. The required storage at each vertex is $\tilde{O}(\frac{1}{\varepsilon}|L_j|/q^i + q^i) = \tilde{O}(\frac{1}{\varepsilon}q^\ell)$.
- At u we store for each $j \in \{1,\ldots,q^i\}$ a vertex $w \in B_i(u)$ with $c_i(w) = j$. By Lemma 6 such a vertex must exist. The storage required for this at u is $\tilde{O}(q^i)$.

In the label of $v \in V$ we store v and $\{p_{L_i}(v), \alpha_i(p_{L_i}(v)), d(v,p_{L_i}(v)), (p_{L_i}(v), v_i')\}_{i=1}^{\ell}$, where $(p_{L_i}(v), v_i')$ is the first edge on a shortest path from $p_{L_i}(v)$ to v.

Routing: A message is routed from u to v as follows. As before, at u we check if there exists $i \in \{0,\ldots,\ell\}$ such that $B_i(u) \cap B_{L_{\ell-i}}(v) \neq \emptyset$ and if this is the case we do the same as in Theorem 13.

Consider the case that for every $i \in \{0,\ldots,\ell\}$ we have $B_i(u) \cap B_{L_{\ell-i}}(v) = \emptyset$. Let $a_i = r_u(\tilde{q}^i)$ and let $b_i = d(v,p_{L_i}(v)) - 1$ if $v \notin L_i$ and $b_i = 0$, otherwise. Let $j \in \arg\min_{i\in\{1,\ldots,\ell\}}\{a_i + b_{\ell-i+1}\}$. If there is more than one index that achieves the minimum we take the index of the highest value. Let $k = \ell - j + 1$. We obtain $\alpha_k(p_k(v))$ from the label of v and route on a shortest path to the vertex $w \in B_j(u)$ with $c_j(w) = \alpha_k(p_{L_k}(v))$, that is stored at u. From w the message is routed to $p_{L_k}(v)$ using Lemma 8. From

$p_{L_k}(v)$ the message is forwarded to v'_k and then using the tree routing scheme of $T_{C_{L_k}}(v'_k)$ it is routed to v. The total length of the path is at most $d(u,w) + (1+\varepsilon)d(w, p_{L_k}(v)) + d(p_{L_k}(v), v)$. Let $d(u,v) = \Delta$.

We first consider the degenerated case of $\Delta = 1$. Since there is no intersection, we have $B_\ell(u) \cap B_{L_0}(v) = \emptyset$ and thereby $a_\ell = 0$. Also we have that $B_{\ell-1}(u) \cap B_{L_1}(v) = \emptyset$, so it follows that $v \notin C_{L_1}(u)$ and $b_1 = 0$. We get that $a_\ell = b_1 = 0$ and by the selection rule of j it must be that $j = \ell$. The message is routed first to $w \in B_\ell(u)$. From there to $p_{L_1}(v)$ and then to v. The length of this path is at most $5 + \varepsilon$.

Consider now the case that $\Delta > 1$. From the triangle inequality it follows that $d(w, p_{L_k}(v)) \leq d(u,w) + d(u,v) + d(v, p_{L_k}(v))$. Therefore, we need to bound $(2+\varepsilon)(a_j + b_k + 2) + (1+\varepsilon)\Delta$.

For every $i \in \{0, \ldots, \ell\}$ we have $B_i(u) \cap B_{L_{\ell-i}}(v) = \emptyset$ thus $a_i + b_{\ell-i} \leq \Delta - 1$. Let $x_0 = a_0/\Delta$ and $x_i = (a_i + 1)/\Delta$, for $i > 0$. Let $y_i = b_i/\Delta$. Since $x_0 = y_0 = 0$ and $x_i + y_{\ell-i} \leq 1$ we can apply Lemma 14 with the series $\{x_i\}_{i=0}^{\ell}$ and $\{y_i\}_{i=0}^{\ell}$ and get $a_j + 1 + b_{\ell-j-1} = a_j + 1 + b_k \leq (1 + 1/\ell)\Delta$. Therefore,

$$\begin{aligned}(2+\varepsilon)(a_j + b_k + 2) + (1+\varepsilon)\Delta &\leq (2+\varepsilon)((1+1/\ell)\Delta + 1) + (1+\varepsilon)\Delta \\ &= \Delta((1+\varepsilon) + (1+1/\ell)(2+\varepsilon)) + 2 + \varepsilon \\ &= \Delta(3 + 2\varepsilon + (2+\varepsilon)/\ell) + 2 + \varepsilon \\ &\leq \Delta(3 + 2/\ell + 4\varepsilon) + 2.\end{aligned}$$

□

Stretch $4k - 7 + \varepsilon$. We present a $(4k-7+\varepsilon)$-stretch routing scheme that uses $O(k \log^2 n / \log\log n)$-bit labels size, $O(\frac{1}{\varepsilon}\log D + \log^2 n/\log\log n)$-bit headers and routing tables of $\tilde{O}(\frac{1}{\varepsilon} n^{1/k} \log D)$ size, where D is the normalized diameter of the graph. This should be compared with the $(4k-5)$-stretch routing scheme of Thorup and Zwick [21] that has $O(k \log^2 n/\log\log n)$-bit labels, $O(\log^2 n/\log\log n)$-bit headers and routing tables of $\tilde{O}(n^{1/k})$ size.

We start with a short overview of the $(4k-5)$-stretch routing scheme of Thorup and Zwick [21]. Let $k \geq 1$ and let $A_0 \supseteq A_1 \supseteq \ldots \supseteq A_k$ be vertex sets, such that $A_0 = V$, $A_k = \emptyset$ and A_i is formed by picking each vertex of A_{i-1} independently with probability $n^{-1/k}$. The expected size of A_i is $n^{1-i/k}$. For every $u \in V$, let $p_i(u) = p_{A_i}(u)$. The *bunch* of $u \in V$ is

$$B(u) = \cup_{i=0}^{k-1} \{v \in A_i \setminus A_{i+1} \mid d(u,v) < d(u, p_{i+1}(u))\},$$

where $d(u, p_k(u)) = \infty$.

The size of every bunch is $O(kn^{1/k})$ in expectation or $\tilde{O}(kn^{1/k})$ in the worst case. The cluster of $w \in A_i \setminus A_{i+1}$ is $C(w) = \{u \in V \mid d(u,w) < d(u, p_{i+1}(u))\}$.

In the preprocessing we compute $\{A_i\}_{i=0}^{k-1}$. Let $w \in V$ and let $T(w)$ be a shortest path tree rooted at w that spans the vertices of $C(w)$. For every $v \in V$, store $B(v)$ at v and the label of v in the tree routing scheme of $T(w)$ for every $w \in B(v)$. The label of v contains $\{p_i(v)\}_{i=0}^{k-1}$ and the tree labels of v in $\{T(p_i(v))\}_{i=0}^{k-1}$.

A message is routed from u to v as follows. At u we find the smallest i for which $u \in C(p_i(v))$. The message is routed using the tree routing scheme of $T(p_i(v))$ from u to v on a path of length at most $d(u, p_i(v)) + d(p_i(v), v)$. Using the triangle inequality we get $d(u, p_i(v)) + d(p_i(v), v) \leq 2d(p_i(v), v) + d(u,v)$. In [21] they show that for every $j \leq i$ it holds that $d(v, p_j(v)) \leq 2d(u,v) + d(v, p_{j-1}(v))$, and therefore, $d(v, p_j(v)) \leq 2jd(u,v)$. Since $j \leq i \leq k-1$ we get that $2d(p_i(v), v) + d(u,v) \leq 2(2(k-1)d(u,v)) + d(u,v) = (4k-3)d(u,v)$.

In [21] they further reduced the stretch to $4k-5$. Notice that $d(v, p_j(v)) \leq 2(j-1)d(u,v) + d(v, p_1(v))$. Using Lemma 4 they compute a set A_1 such that $|C(u)| = O(n^{1/k})$ for every $u \in A_0 \setminus A_1$. Each such u

stores the set $C(u)$ and for every $v \in |C(u)|$ the label of v in the tree routing scheme of $T(u)$. To route a message from $u \in A_0 \setminus A_1$ to v, u checks if $v \in C(u)$ and if it is then the message is routed to v on $T(u)$. If not, then it follows that $d(v, p_1(v)) \leq d(u, v)$ (as opposed to $d(v, p_1(v)) \leq 2d(u, v)$ before) and the stretch is $4k - 5$. See [21] for the full description.

We now turn to present our improved routing scheme.

Theorem 16. *Let $G = (V, E, \omega)$ be a weighted undirected graph, where $\omega : E \to \mathbb{R}^+$. For every $\varepsilon > 0$, there is a routing scheme for G that uses $o(k \log^2 n)$-bit labels, $\tilde{O}(\frac{1}{\varepsilon} n^{1/k} \log D)$ space at each vertex and $\tilde{O}(\frac{1}{\varepsilon} \log D)$-bit header, to route a message between any pair of vertices on a $(4k - 7 + \varepsilon)$-stretch path.*

Proof.
Preprocessing: We store the same information as in the $(4k - 5)$-stretch routing scheme of [21].

Let $q = n^{1/k}$. Additionally, we store $B(u, \tilde{q})$ at u, for every $u \in V$. This requires $\tilde{O}(n^{1/k})$ space at each vertex.

Let c be a coloring function with q colors obtained using Lemma 6 with respect to the sets $B(u, \tilde{q})$, for every $u \in V$. Let $\mathcal{U} = \{U_1, \ldots, U_q\}$ be the color sets induced by c. Let $\mathcal{W} = \{W_1, \ldots, W_q\}$ be an arbitrary partition of A_{k-2} into q sets each with at most $|A_{k-2}|/q = \tilde{O}(n^{2/k}/n^{1/k}) = \tilde{O}(n^{1/k})$ vertices. For $w \in A$, let $\alpha(w)$ be the index of the set in partition $\mathcal{W}$ that contains w, that is $w \in W_{\alpha(w)} \in \mathcal{W}$. We use Lemma 8 with the partitions $\mathcal{U}$ and $\mathcal{W}$. The required storage at each vertex is $\tilde{O}((\frac{1}{\varepsilon} \log D)|A_{k-2}|/q + q) = \tilde{O}(\frac{1}{\varepsilon} n^{1/k} \log D)$.

It follows from Lemma 6 that $B(u, \tilde{q})$ contains a vertex of each color. For every $i \in \{1, \ldots, q\}$ we store at u a vertex from $B(u, \tilde{q})$ of color i. The storage required for this at u is $\tilde{O}(n^{1/k})$.

The label of every $v \in V$ is composed of the label of v in the $(4k-5)$-stretch routing scheme and $\alpha(p_{k-2}(v))$.

Routing: A message is routed from u to v as follows. At u we check if $v \in B(u, \tilde{q})$. If this is the case the message is routed from u to v on a shortest path. If this is not the case we look for the smallest index i for which $u \in C(p_i(v))$. If $i \leq k - 2$, the message is routed to v using the tree routing scheme of $T(p_i(v))$ on a path of length at most $d(u, p_i(v)) + d(p_i(v), v)$. It follows from the analysis of [21] that $d(u, p_i(v)) + d(p_i(v), v) \leq (4k - 9)d(u, v)$, for $i \leq k - 2$. If $i = k - 1$ we route the message to a vertex $w \in B(u, \tilde{q})$ for which $c(w) = \alpha(p_{k-2}(v))$. The message is routed from w to $p_{k-2}(v)$ using Lemma 8, and from $p_{k-2}(v)$ to v using the tree routing scheme of $T(p_{k-2}(v))$ (recall that the label of v in the tree routing scheme of $T(p_{k-2}(v))$ is in v's label). The length of the path is at most $d(u, w) + (1+\varepsilon)d(w, p_{k-2}(v)) + d(p_{k-2}(v), v)$. Using the triangle inequality we can bound it with $d(u, v) + (1+\varepsilon)(2d(u, v) + d(p_{k-2}(v), v)) + d(p_{k-2}(v), v)$. From the analysis of [21] it follows that $d(p_{k-2}(v), v) \leq (2(k - 3) + 1)d(u, v)$ and we get that the stretch is $4k - 7 + \varepsilon$. □

References

[1] I. Abraham and C. Gavoille. On approximate distance labels and routing schemes with affine stretch. In David Peleg, editor, *DISC*, volume 6950 of *Lecture Notes in Computer Science*, pages 404–415. Springer, 2011.

[2] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. In Phillip B. Gibbons and Micah Adler, editors, *SPAA*, pages 20–24. ACM, 2004.

[3] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. *ACM Transactions on Algorithms*, 4(3), 2008.

[4] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999.

[5] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

[6] M. Arias, L. Cowen, K. A. Laing, R. Rajaraman, and O. Taka. Compact routing with name independence. *SIAM J. Discrete Math.*, 20(3):705–726, 2006.

[7] B. Awerbuch, A. Bar-Noy, N. Linial, and D. Peleg. Compact distributed data structures for adaptive routing (extended abstract). In *STOC*, pages 479–489, 1989.

[8] B. Awerbuch, A. Bar-Noy, N. Linial, and D. Peleg. Improved routing strategies with succinct tables. *J. Algorithms*, 11(3):307–341, 1990.

[9] S. Baswana and S. Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007.

[10] S. Chechik. Compact routing schemes with improved stretch. In Panagiota Fatourou and Gadi Taubenfeld, editors, *PODC*, pages 33–41. ACM, 2013.

[11] L. Cowen. Compact routing with minimum stretch. *J. Algorithms*, 38(1):170–183, 2001.

[12] D. Dor, S. Halperin, and U. Zwick. All-pairs almost shortest paths. *SIAM J. Comput.*, 29(5):1740–1759, 2000.

[13] T. Eilam, C. Gavoille, and D. Peleg. Compact routing schemes with low stretch factor. *J. Algorithms*, 46(2):97–114, 2003.

[14] P. Erdős. Extremal problems in graph theory. *Theory of graphs and its applications*, pages 29–36, 1964.

[15] P. Fraigniaud and C. Gavoille. Routing in trees. In Fernando Orejas, Paul G. Spirakis, and Jan van Leeuwen, editors, *ICALP*, volume 2076 of *Lecture Notes in Computer Science*, pages 757–772. Springer, 2001.

[16] D. Peleg and A. A. Schäffer. Graph spanners. *J. Graph Theory*, pages 99–116, 1989.

[17] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Computing*, 18(4):740–747, 1989.

[18] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.

[19] M. Pătraşcu and L. Roditty. Distance oracles beyond the Thorup-Zwick bound. *SIAM J. Comput.*, 43(1):300–311, 2014.

[20] M. Pătraşcu, L. Roditty, and M. Thorup. A new infinity of distance oracles for sparse graphs. In *FOCS*, pages 738–747, 2012.

[21] M. Thorup and U. Zwick. Compact routing schemes. In *SPAA*, pages 1–10, 2001.

[22] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005.